Experimental Evaluation of Centralized Failure Restoration in a Dynamic Impairment-Aware All-Optical Network

J. Perelló[1], S. Spadaro[1], F. Agraz[1], M. Angelou[1,2], S. Azodolmolky[1,2], Y. Qin[3], R. Nejabati[3], D. Simeonidou[3], P. Kokkinos[4], E. Varvarigos[4], S. Al Zahr[5], M. Gagnaire[5], and I. Tomkos[2]

(1*) Universitat Politècnica de Catalunya (UPC), Barcelona, Spain; (2) Athens Information Technology (AIT), Athens, Greece; (3) University of Essex, Colchester, United Kingdom; (4) Research Academic Computer Technology Institute (RACTI), Patras, Greece; (5) Telecom ParisTech (ENST), Paris, France; e-mail address: perello@ac.upc.edu*

Abstract: A centralized impairment-aware lightpath restoration scheme is experimentally demonstrated. Through the implementation of the QoT estimator module on FPGA technology, lightpath restoration times around 1.3s are obtained for the high-priority traffic class.

OCIS codes: (060.4261) Networks, protection and restoration

1. Introduction

Next-generation optical transport networks are expected to evolve from those nowadays' static opaque networks to dynamic all-optical ones, where data transmission and switching are performed entirely in the optical domain and everything is orchestrated from a distributed control plane entity. Main drivers to this migration can be found in the reduced network costs resulting from the end-to-end optical transparency, which relieves operators from deploying service-dependent intermediate electronic processing stages, together with the automated control-plane-driven connection (i.e., *lightpath*) provisioning and restoration.

Nonetheless, the realization of such dynamic all optical networks has not yet been achieved for commercial exploitation so far. Unfortunately, the end-to-end transparency may lead to an unfeasible communication due to the physical layer degradations that the signal accumulates along the path. Moreover, such a transparency complicates failure localization and isolation procedures, as Loss-of-Light (LoL) alarms propagate along the path. In this context, the EU FP7 DICONET project [1] has addressed both challenges looking towards the future Internet backbone. The main outcome of the project has been the development of a Network Planning and Operation Tool (NPOT) that gathers real-time optical layer performance metrics and incorporates them into Impairment-Aware Routing and Wavelength Assignment (IA-RWA) algorithms. These performance metrics are delivered to the NPOT through an impairment-aware extended Generalized Multi-Protocol Label Switching (GMPLS) control plane. Besides, a built-in failure localization module is placed in the NPOT enabling a successful failure recovery. In order to ensure fast route computations during the network operation phase, the Quality of Transmission (QoT) estimator in the NPOT (hereafter the Q-Tool) is implemented on Field Programmable Gate Array (FPGA) hardware.

This work reports the experimental evaluation of a centralized impairment-aware failure restoration in the 14-Node DICONET test-bed. The obtained results are compared to the previous work in [2], where the failure localization functionality was left to the GMPLS Link Management Protocol (LMP) and a non-accelerated (i.e., software-based) Q-Tool was used. As an additional contribution of this work, a prioritized NPOT scheduler has been implemented so that those lightpath restorations with higher priority can be served first and, thus, they can experience lower restoration times.

2. Impairment-aware centralized lightpath restoration scheme

The impairment-aware lightpath restoration scheme devised in DICONET relies on a centralized NPOT common to all network entities, as depicted in Fig. 1. The goal of the NPOT is the computation of valid routes upon lightpath establishment/restoration request. To achieve this, the NPOT maintains two different databases, namely, the global Traffic Engineering Database (gTED) and the global Physical Parameter Database (gPPD), which store the current wavelength availability and Physical Layer Impairment (PLI) information in the network, respectively. These databases are updated through an extended version of the GMPLS OSPF-TE protocol, allowing the inclusion and dissemination of the links' PLI values in the Opaque Link State Advertisements (OLSAs).

The main module in the NPOT is the Q-Tool, able to estimate, taking both linear (ASE, CD, FC and PMD) and non-linear PLIs (SPM, XPM, FWM) into account, the feasibility (in terms of the Q-factor value) of a lightpath to be established in the network, in addition to that of the potentially affected active ones. This is a computationally intensive task, though, arising as the route computation bottleneck. Hence, the Q-Tool module in the DICONET centralized NPOT has been implemented on FPGA hardware and coupled to the NPOT following a client-server model [3]. Specifically, the Q-Tool is composed of two different modules: the QoT Estimation Agent (QoTEA) and Estimation Engine (QoTEE).

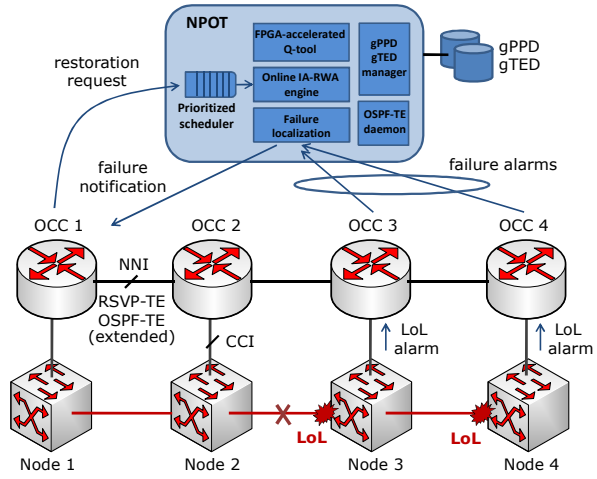

Fig. 1. Centralized restoration example: a bidirectional lightpath is established from node 2 to 4 and a failure occurs between nodes 2 and 3.

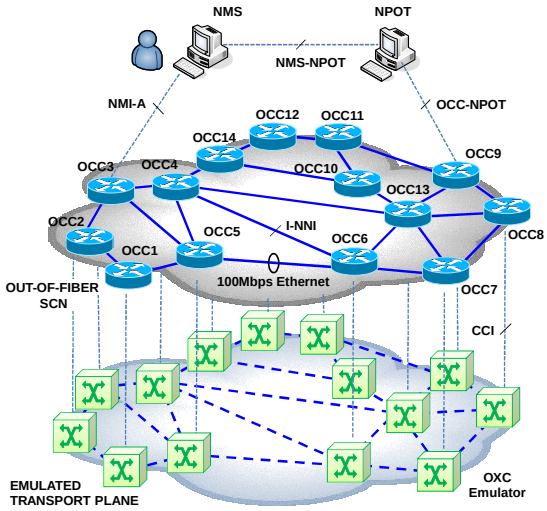

Fig. 2. The DICONET test-bed used for the experiments: A 14-node meshed transport plane with 23 links is emulated.

The QoTEE module is deployed in a Xilinx Virtex IV FPGA and is the responsible for the actual QoT estimation. In turn, the QoTEA runs on a 300MHz IBM PowerPC 405 hard core embedded inside the FPGA fabric with 1GB DDR2 memory, and is responsible for receiving the lightpath QoT estimation requests and sending them to the QoTEE. These requests come from the online IA-RWA engine, which computes feasible routes for the new traffic demands taking the current network state into account (i.e., using the information in both gPPD and gTED databases). For the specific details about the implemented IA-RWA algorithm please refer to [4].

Route requests are served on a one-by-one basis in the current centralized NPOT version. Having this in mind, a scheduler has been placed in the NPOT, which stores the new incoming requests until the online IA-RWA engine becomes available. As will be illustrated in the experimental results, this scheduler plays an important role in the failure restoration process, where a significant amount of restorations must be handled almost simultaneously. Targeting at low restoration time figures for the high-priority restorable traffic, two priorities have been introduced there. In this way, high-priority restoration requests can overtake the low-priority ones, thus being served first. For all and all, an efficient failure recovery would not be possible without the failure localization module. This module is initially employed in the network planning phase to design an *m-trail* solution able to localize the broken link in the network with low monitoring deployment CAPEX [5]. Basically, a different code is assigned to each monitor in the network. Hence, based on the LoL alarms arising from a failure, which contain the specific monitor code, and the alarm code table built during the *m-trail* design, the module can succeed in localizing the exact failed link.

Fig. 1 illustrates an example of the implemented centralized restoration scheme, where a failure affecting the bidirectional lightpath from nodes 1 to 4 is assumed. Provided that the downstream nodes 3 and 4 are equipped with Optical Performance Monitors (OPMs), they detect LoL in their incoming ports, which is notified to the respective Optical Connection Controllers (OCCs). Such OCCs detecting the failure also inform the centralized NPOT about the failure state and the code of the OPM that has detected it. By doing so, the failure localization module can localize and isolate the failure, updating gPPD and gTED databases accordingly. Furthermore, it can also notify the failure to the source nodes of the affected restorable lightpaths, so that they can start the restoration procedures. In the implemented centralized restoration scheme, failure restoration is delegated to the GMPLS control plane, in particular to the standard RSVP-TE protocol. Therefore, as soon as the source node is notified about the failure, it requests a backup route to the centralized NPOT for restoration purposes. Being this route available, it is returned to the source node, which triggers the RSVP-TE protocol to establish the backup path.

3. Experimental evaluation

The performance of the proposed centralized lightpath restoration approach has been validated on the DICONET test-bed, located at the UPC premises (Fig. 2). The test-bed describes the same topology as the 14-node Deutsche Telekom (DT) network [2], where 10 bidirectional wavelengths per link have been assumed. Each network node is composed of an OCC and a WSS-based OXC emulator interconnected through the CCI interface. In turn, all OCCs are interconnected to the NMS and the NPOT through the NMI-A and OCC-NPOT interfaces, respectively.

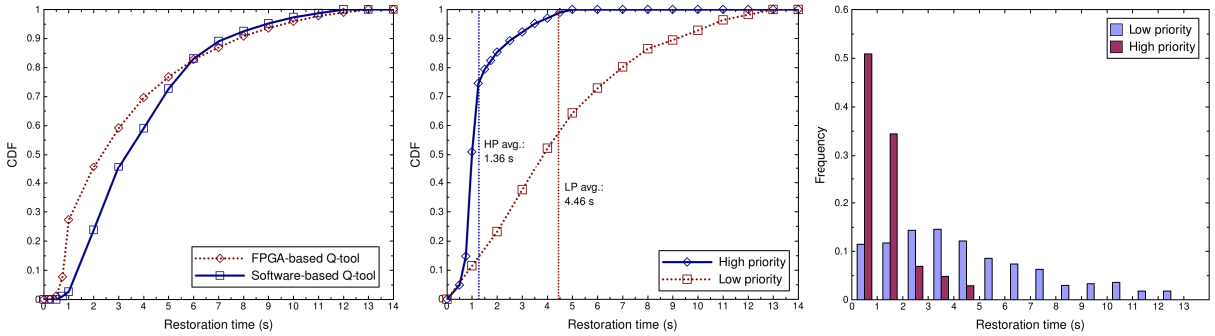

Fig. 3. Performance evaluation: lightpath restoration time CDF with and without FPGA acceleration (left); lightpath restoration time CDF for high and low priority traffic classes (center); lightpath restoration time frequency for high and low priority traffic classes (right).

The connectivity between OCCs is supported over 100 Mbps point-to-point Ethernet links, which describe an out-of-fiber control plane with the same topology as the emulated all-optical data plane. OCCs implement the whole GMPLS protocol set. As in [2], we initially load the network with 10, 20, 30, 40 and 50 bidirectional lightpaths between randomly selected node pairs. These lightpaths can be either 1+1 protected, directly establishing QoT-compliant working and backup paths for them, or restorable, for which a backup path establishment is triggered in case of working path failure. The same 70-30% restorable-protected ratio used in [2] has been here assumed. Besides, in order to obtain the results hereafter presented, 10 independent failures are generated in each of the 5 network scenarios. Initially, no priorities are set in the NPOT scheduler.

Fig. 3 (left) presents the Cumulative Distribution Function (CDF) of the lightpath restoration time in the network. Moreover, the same results when a software-based Q-Tool is employed are also plotted as a benchmark. As shown, although the FPGA-accelerated Q-Tool leads to substantial restoration time befits (e.g., around 30% of the restorations are performed below 1s), the last requests to arrive still suffer from high queuing delays due to the sequential NPOT behavior. As a matter of fact, a guard time is also left between two consecutive route computations so that the NPOT is fed with the new wavelength availability and PLI information (around 2s in the test-bed). One may argue that waiting for the OSPF-TE flooding completion between route computations is quite conservative. Nonetheless, it really serves the purpose of minimizing the network blocking probability.

Aiming to provide a differentiated service to those restorable lightpaths with higher priority, two different priorities (high and low priority) have been defined in the NPOT scheduler, as explained previously in section 2. In particular, a 40-60% high-low priority ratio has been assumed in the conducted experimental results. Fig. 3 (center) depicts the CDF of the high and low priority lightpaths' restoration time in the network. As observed, by implementing differentiated queuing in the NPOT scheduler, attractive restoration times around 1.36s in average can be provided to the high priority restorable traffic. For the low priority traffic, however, the average restoration time increases from 3.22 to 4.46s compared to the scenario without priorities, as the low priority restoration requests are the last ones to be processed. This would not be really significant, though, if we assign this class of service to the best-effort traffic supported on the network. For better illustration, the lightpath restoration time frequency is also plotted in Fig. 3 (right). As seen, more than half of the high priority restorations are concentrated on the interval between 0 and 1s (i.e., sub-second restoration time). Furthermore, none of them experiences a restoration time higher than 5s. In contrast, the low priority traffic shows a longer tail, until 13s, having the peak between 3 and 4s.

4. Conclusions

This paper reported the evaluation of the centralized impairment-aware lightpath restoration proposed and validated within the EU DICONET Project. This work has been supported by the European Commission through the FP7 DICONET Project and the Spanish Science Ministry through the project ENGINE (TEC2008-02634).

References

[1] S. Azodolmolky, et al., "A dynamic impairment aware networking solution for transparent meshed optical networks", IEEE Communications Magazine, vol. 47, no. 5, pp. 38-47, May 2009.
[2] F. Agraz et al., "Experimental Evaluation of Path Restoration for a Centralized Impairment-Aware GMPLS-controlled All-Optical Network", in Proc. of 36th European Conference and Exhibition on Optical Communication (ECOC 2010), Torino, Italy, September 2010.
[3] Y. Qin et al., "Demonstration of C/S based Hardware Accelerated QoT Estimation Tool in Dynamic Impairment-Aware Optical Network", in Proc. of 36th European Conference and Exhibition on Optical Communication (ECOC 2010), Torino, Italy, September 2010.
[4] P. Kokkinos, K. Christodoulopoulos, K. Manousakis, and E. Varvarigos, "Multi-parametric online RWA based on impairment generating sources," in Proc. of IEEE Global Communications Conference (GLOBECOM 2009), Honolulu, Hawaii, USA, November 2009.
[5] M. Gaugnaire, et al., "Failure Localization Algorithms to Enable Resiliency," DICONET Deliverable D4.4.